Accelerated Design and Deployment of Low-Carbon Concrete for Data Centers

Xiou Ge*
Ming Hsieh Department of Electrical and Computer Engineering, University of Southern California
USA

Richard T. Goodwin
IBM Thomas J. Watson Research Center
USA

Haizi Yu
Knowledge Lab, University of Chicago
USA

Pablo Romero
Department of Civil and Environmental Engineering, University of Illinois Urbana-Champaign
USA

Omar Abdelrahman
Department of Civil and Environmental Engineering, University of Illinois Urbana-Champaign
USA

Amruta Sudhalkar
Meta
USA

Julius Kusuma
Meta
USA

Ryan Cialdella
Ozinga Ready Mix
USA

Nishant Garg
Department of Civil and Environmental Engineering, University of Illinois Urbana-Champaign
USA

Lav R. Varshney
Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign
USA

Figure 1: Depiction of Meta's data center in DeKalb, IL, USA where low-carbon concrete discovered by artificial intelligence methods has been tested.

*This work was performed when the author was at the University of Illinois Urbana-Champaign.

COMPASS '22, June 29–July 1, 2022, Seattle, WA, USA

ABSTRACT

Concrete is the most widely used engineered material in the world with more than 10 billion tons produced annually. Unfortunately, with that scale comes a significant burden in terms of energy, water, and release of greenhouse gases and other pollutants; indeed 8% of worldwide carbon emissions are attributed to the production

ACM ISBN 978-1-4503-9347-8/22/06…$15.00
https://doi.org/10.1145/3530190.3534817

of cement, a key ingredient in concrete. As such, there is interest in creating concrete formulas that minimize this environmental burden, while satisfying engineering performance requirements including compressive strength. Specifically for computing, concrete is a major ingredient in the construction of data centers.

In this work, we use conditional variational autoencoders (CVAEs), a type of semi-supervised generative artificial intelligence (AI) model, to discover concrete formulas with desired properties. Our model is trained just using a small open dataset from the UCI Machine Learning Repository joined with environmental impact data from standard lifecycle analysis. Computational predictions demonstrate CVAEs can design concrete formulas with much lower carbon requirements than existing formulations while meeting design requirements. Next we report laboratory-based compressive strength experiments for five AI-generated formulations, which demonstrate that the formulations exceed design requirements. The resulting formulations were then used by Ozinga Ready Mix—a concrete supplier—to generate field-ready concrete formulations, based on local conditions and their expertise in concrete design. Finally, we report on how these formulations were used in the construction of buildings and structures in a Meta data center in DeKalb, IL, USA. Results from field experiments as part of this real-world deployment corroborate the efficacy of AI-generated low-carbon concrete mixes.

CCS CONCEPTS

• **Applied computing** → **Engineering**; • **Computing methodologies** → **Machine learning**.

KEYWORDS

sustainable building materials, artificial intelligence, concrete, variational autoencoders

ACM Reference Format:
Xiou Ge, Richard T. Goodwin, Haizi Yu, Pablo Romero, Omar Abdelrahman, Amruta Sudhalkar, Julius Kusuma, Ryan Cialdella, Nishant Garg, and Lav R. Varshney. 2022. Accelerated Design and Deployment of Low-Carbon Concrete for Data Centers. In *ACM SIGCAS/SIGCHI Conference on Computing and Sustainable Societies (COMPASS '22), June 29–July 1, 2022, Seattle, WA, USA.* ACM, New York, NY, USA, **??** pages. https://doi.org/10.1145/3530190.3534817

1 INTRODUCTION

Is it possible to formulate concrete mixtures with standard ingredients so that they have half as much embodied carbon as traditional formulations, yet are just as strong? We answer this question affirmatively, through generative artificial intelligence (AI) algorithms, lab testing, and field trials at Meta's data center in DeKalb, IL, USA.

The building sector accounts for a significant fraction of overall energy consumption and pollution worldwide, including residential, commercial, and infrastructural structures such as the construction of computing infrastructure like data centers. Concrete—including its key ingredient, cement—is one of the most energy-intensive and polluting building materials to fabricate, but is also the most widely used engineered material in the world with more than 10 billion tons produced annually. Unfortunately, with that scale comes a significant burden in terms of energy, water, and release of greenhouse gases and other pollutants; indeed 8% of worldwide carbon emissions are attributed to the production of cement [2, 10]. As part of larger efforts to use machine learning to address problems in climate change [33], here we approach the design of low-carbon concrete using generative machine learning (ML) algorithms. But we do not just stop at the algorithm stage; rather we further demonstrate the efficacy of our approach via laboratory experiments, translation to industrial practice, and field tests in a real data center location, which confirms the practicality of our approach. Figure 2 depicts how we move from stage to stage of our research.

Although concrete has been designed, developed, and optimized for more than 7000 years across numerous civilizations, yielding famous structures such as the Pantheon in Rome, the Willis Tower in Chicago, and the Burj Khalifa in Dubai [5, 17, 29], there are still numerous open questions regarding its properties. As a material, concrete is primarily composed of water, fine/coarse aggregate, cement, and other cementitious materials such as slag and fly ash, which are industrial byproducts. Slag is a stony material that is waste from steel production, whereas fly ash is composed of particulates that remain from coal combustion. Only very recent work in computational statistical mechanics and theory has started to reveal how cement cohesion happens at the nanoscale level to give concrete its strength over time [14], but physics-based models are unable to accurately predict mechanical properties of concrete. As such, there is a growing scientific literature on using supervised machine learning to predict concrete strength based on its composition [4, 41, 42]. Most such extant ML predictors draw on ensemble methods such as random forests, which have limited scientific interpretability [3] and so it is unclear how to directly use insights from predictive algorithms to design novel concrete formulations with particular desired properties.

Since the first presentation of our work in 2019 on generating novel concrete formulations [13]—which was restricted to just the algorithmic stage—there have been several papers that put multi-objective optimization procedures around predictive ML for design of concrete formulations. All resort to meta-heuristics for global optimization [6, 8, 16, 26, 32, 35, 43, 44], an approach that is not only very inefficient computationally, but may also use predictive ML algorithms well-outside the domain of their training data where their accuracy is quite limited [24, 28]. These recent works do not perform laboratory or field experiments to validate the results of the algorithms. We take an alternative approach and develop *generative* ML models rather than predictive ones: such models directly generate new concrete formulations rather than requiring an outer optimization loop and make as best use of training data as they can.

Indeed, recent advances in machine learning and artificial intelligence have enabled machines to generate very high-quality artifacts, such as images of realistic looking faces with certain desired properties [40], or natural language with certain desired topics and styles [19]. In this work, we use an extension of the popular generative modeling framework of variational autoencoders (VAEs) [22] that allows control of attributes. This extension is conditional variational autoencoders (CVAEs) [34], a type of semi-supervised generative model [21], which we can use to generate concrete formulas with desired properties. Note that although deep generative

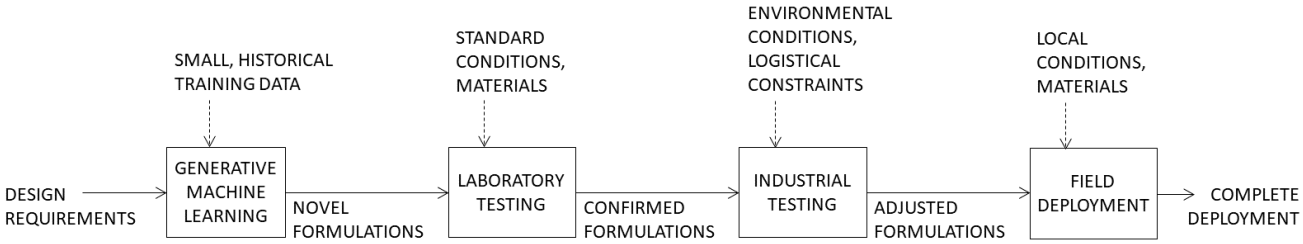

Figure 2: Schematic of the accelerated design and deployment process we undertook, going from training a generative algorithm to full deployment in the field.

models have been applied in materials and molecules discovery—[15] uses VAEs based on recurrent neural networks for chemical design in which molecules are encoded as strings and [31] uses VAEs to improve the accuracy of drug response predictions—they have not previously been used in the context of aggregate materials. In fact, the field of ML-based accelerated materials discovery is oriented away from aggregate materials such as concrete [25], e.g. the Materials Genome Initiative [18] is focused on inorganic compounds, nonporous materials, and electrodes.

We demonstrate CVAEs can design concrete formulas with lower emissions and natural resource usage while meeting design requirements, computationally and also in laboratory and field experiments. In the computational phase before experimental testing, we also train regression models to predict the environmental impacts and strength of generated formulas. This provides initial insight to civil engineers in creating formulas that meet structural needs and best addresses local environmental concerns.

To assess environmental impacts, the Cement Sustainability Initiative (CSI) developed the Environmental Product Declaration (EPD) tool to facilitate the generation of sector-specific EPDs for cement, concrete, clinker, lime, and plaster. EPD is a voluntary declaration that provides quantitative information about the environmental impact of a product, using life-cycle assessment (LCA) methodology and verified by an independent third party. The cloud-based tool was designed to be easy-to-use, to facilitate the process overall, and to reduce the costs of preparing cement and concrete EPDs. In this work, we join this with open, historical concrete formulation and strength data from the UCI Machine Learning Repository [41].

1.1 Sustainable Data Centers

Much of the conversation on sustainable computing infrastructure, such as data centers, focuses on operational sustainability [12, 27] by ensuring facilities use as little power as possible, use renewable energy sources, and have minimal impacts on local areas via co-generation and related techniques. Yet, besides communication, computation, and storage of digital information that operational efficiency captures, there is also the expressly material infrastructure in buildings and structures [11]. This is also starting to get attention [23], with a particular focus on embodied carbon in building infrastructure such as concrete. Embodied carbon includes emissions caused by extraction, manufacture, transportation, assembly, maintenance, replacement, deconstruction, disposal, and end of life

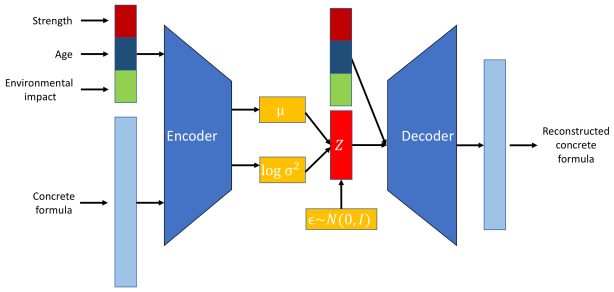

Figure 3: CVAE model structure

aspects of the materials and systems that make up a building. As such, we conduct our field deployment in the context of a new data center being built in DeKalb, IL, USA, see Figure 1.

1.2 Paper Organization

The remainder of the paper is organized as follows. We first describe the data set and the CVAE model details. Next we give results, first showing the average percentage reduction environmental impact achieved by generated better-performing concrete formulas. We then show strength spectrum plots in the 3D environmental impact space which could be turned into a visualization tool for concrete designers. Third, we evaluate the performance of strength conditioned generation of the trained model. After detailing the computational phase of our research, we then present results from laboratory experiments and from full field deployments.

2 GENERATIVE MACHINE LEARNING

2.1 Training Data

We train our model using the Concrete Compressive Strength Data Set [41] openly available from the UCI Machine Learning Repository. It has 1,030 training examples, with seven continuous features describing the amount of constituent material such as cement, aggregates, and water. Compressive strength, after a particular curing time (age), of each concrete formula is also given. In addition, we use the CSI EPD tool to estimate the environmental impact of each concrete formula. The EPD tool produces 12 continuous features characterizing the concrete environmental impact. Among these, we largely focus on global warming potential (GWP) measured using embodied CO_2, acidification potential (AP) measured using

embodied SO_2, and concrete batching water (CBW) consumption. Future work could more explicitly focus on the other nine environmental impacts.[1]

2.2 Generative Model

Our model is based on a variant of the VAE [22] called CVAE [34] as shown in Figure 3. Like other generative models, the goal is to estimate the data distribution $p(y)$ and to generate realistic new samples from that distribution [9]. What makes CVAE different from VAE is that instead of merely generating realistic samples from the data distribution $p(y)$ randomly, we generate from the conditional distribution $p(y|x)$ which give us control over the underlying properties of generated data by conditioning on different values of x.

We interpret the variables in the conditional generative model as follows: x represents the side information of a formula including the strength, age, and environmental impacts, y represents the constituent material amount of a formula, and z represents the latent variables. Like the VAE, a CVAE consists of an encoder $q_\phi(z|x,y)$ that maps the data points to latent codes and a decoder $p_\theta(y|x,z)$ that reconstructs the data points from latent codes. The decoder and encoder are implemented as neural networks where ϕ and θ are the respective network parameters. Since the goal is to generate realistic concrete samples with desired properties, we want to maximize the log-likelihood of the data distribution model $\log p_\theta(y^{(i)}|x^{(i)})$. Since the data distribution $p_\theta(y|x)$ and the posterior distribution $p_\theta(z|x,y)$ are both intractable, we maximize the Evidence Lower Bound (ELBO), $\mathcal{L}$, instead. The loss function of CVAE is therefore:

$$\log p_\theta(y^{(i)}|x^{(i)}) \geq \mathbf{E}_{z\sim q_\phi(z|x,y)}[\log p_\theta(y^{(i)}|z,x^{(i)})] \qquad (1)$$
$$- D_{\mathrm{KL}}(q_\phi(z|x^{(i)},y^{(i)}) \,||\, p_\theta(z|x^{(i)})) = \mathcal{L},$$

where $D_{\mathrm{KL}}(\cdot\,||\,\cdot)$ is the Kullback-Leibler divergence.

2.3 Implementation Details

In our model, the encoder network consists of four fully-connected layers with 25 neurons on the first layer, 20 neurons on the second layer, followed by two parallel layers with two neurons on each which represent the mean and log variance respectively. The prior is set to be an isotropic Gaussian distribution with zero mean and unit variance $p(z) = \mathcal{N}(0, \mathcal{I})$. The reparameterization trick is performed to make the sampling step differentiable and enable backpropagation for training. The decoder network consists of two fully-connected layers with 20 neurons on the first layer and 25 neurons on the second layer. Rectified linear unit (ReLU) activation functions are applied to all layers except the output layer of the decoder, where we use sigmoid activation since we scale our data to $[0, 1]$. The model is trained end-to-end with a version of stochastic gradient descent: the Adam optimizer with learning rate of 0.001 and batch size of 10 [20].

[1]The other nine metrics are eutrophication potential (EP), formation potential of tropospheric ozone photochemical oxidants (POCP), depletion potential of the stratospheric ozone layer (ODP), total primary energy consumption (PEC), depletion of non-renewable energy (NRE), use of renewable primary energy (RE), depletion of non-renewable material resource (NRM), use of renewable material resources (RM), concrete hazardous waste (CHW), and concrete non-hazardous waste (CNHW).

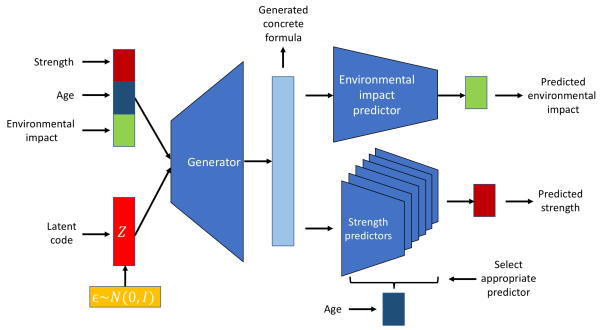

Figure 4: Generating new concrete formulas and evaluating their properties

Metric	GWP (kg CO_2 eq./m^3)	AP (kg SO_2 eq./m^3)	CBW (m^3)
MAE	7.187	0.019	0.003
RMSE	9.374	0.040	0.006
R^2	0.979	0.974	0.881

Table 1: Environmental Impacts Predictor Performance

2.4 Property Predictors

We also trained neural network-based regression models as shown in Figure 4 using the dataset that we described above to predict the environmental impact and strength of concrete formulas. Since the compressive strength is dependent on the age of concrete, we trained separate compressive strength predictors for each age group. The purpose of the predictors is twofold. First, we can measure how well the properties of generated samples match the desired properties given as conditioning variables during data generation. Second, we can make fair comparisons between extant and generated concrete formulas in terms of the environmental impact. We experimented with three different types of regression models, namely linear regression, decision tree regression, and neural network regression. Although linear regression can achieve comparable performance with decision tree regression and neural network regression, it often predicts far-out-of-range values for newly generated concrete formulas. The neural network regression has slightly better performance than the decision tree regression and therefore we use the former for prediction tasks. The performance of the neural network regression models for global warming potential (GWP), acidification potential (AP), and concrete batching water (CBW) consumption are shown in Table 1. The performance of the strength predictors are shown in Table 2, where compressive strength is measured in megapascals (MPa) for several different ages.

3 COMPUTATIONAL RESULTS

3.1 Concrete Formulas with Reduced Environmental Impact

To demonstrate that the generative algorithm discovers new concrete formulas with reduced environmental impacts, we compared

	Predictor Performance (MPa)					
	≤3 days	7 days	14 days	28 days	56 days	≥90 days
MAE	2.985	3.850	3.378	6.015	5.093	4.457
RMSE	0.222	0.201	0.163	0.227	0.124	0.125
R^2	0.819	0.870	0.703	0.679	0.795	0.789

Table 2: Strength Predictor Performance

		Average Reduction (%)		
Age (day)	Strength (MPa)	GWP (kg CO_2 eq./m^3)	AP (kg SO_2 eq./m^3)	CBW (m^3)
≤3	30±1	0.80	1.83	5.47
	40±1	7.74	1.59	0.26
7	30±1	19.69	3.94	7.58
	40±1	25.45	11.33	5.03
14	20±1	2.20	5.72	10.64
	60±1	42.45	21.09	5.17
28	70±1	21.62	6.66	3.32
	80±1	27.44	8.40	4.15
56	40±1	4.38	2.95	7.04
	50±1	14.38	3.23	3.64
	70±1	30.26	23.75	1.32
	80±1	5.88	1.33	3.46
≥90	80±1	30.58	6.91	4.11

Table 3: Average environmental impact reduction achieved by better performing generated samples

Strength (MPa)	30±1	40±1
Constituent Material	Amount (kg per m^3)	
Cement	186.4	259.0
Blast Furnace Slag	236.7	288.6
Fly Ash	107.1	58.8
Water	142.3	142.5
Superplasticizer	22.3	26.1
Coarse Aggregate	901.4	868.6
Fine Aggregate	717.2	763.0

Table 4: Sample concrete formula with reduced environmental impact

the GWP, AP, and CBW values between existing concrete formulas in the training set [41] and generated formulas with the same age and similar strength. For each concrete age group, we generate 60,000 concrete formulas. Both the strength and the environmental impact inputs to the generator are produced by randomly sampling from the standard uniform distribution whereas the latent code input is produced by randomly sampling from the standard bivariate normal distribution. We then use the trained environmental impact

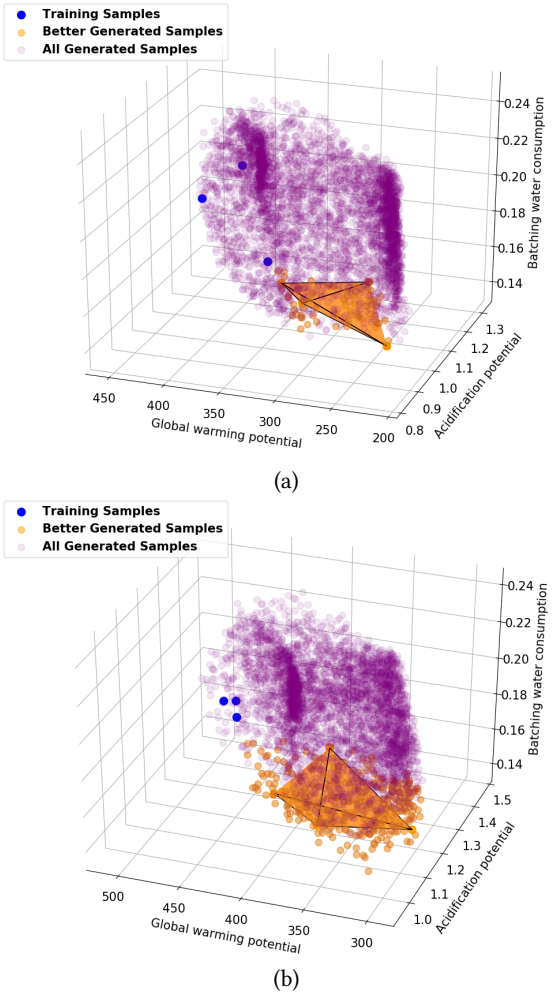

Figure 5: Approximated hull of generated samples from archetypal analysis, training samples, and all generated samples for specific curing time and strength level. (a) Curing time = 7 days, Strength = 30±1 MPa. (b) Curing time = 7 days, Strength = 40±1 MPa.

predictor and strength predictor for the corresponding age group to evaluate environmental impact and strength of the newly generated formulas. We count the number of generated samples having lower environmental impact than the best observed values for extant samples in all 3 dimensions. We also measured the average percentage reduction in environmental impact for the better-performing samples as compared to extant samples in the training set.

Notice that this performance metric of *conditional average improvement* is quite conservative, as it is not just considering the best generated formulations but the whole ensemble of formulations that are better than formulations in the training dataset along all three environmental performance dimensions.

Results in Table 3 show significant improvements in environmental impact. Taking 14-day strength around 60 MPa, notice that the conditional average reduction for carbon (GWP) can be as high

as 42%, while also achieving conditional reduction for sulfur (AP) as high as 21%. As noted, conditional average improvement/reduction are very conservative performance metrics. When we look at some of the best specific formulations that emerge from the generative algorithm, we will see more than 50% reduction in carbon. An example is given in Table 4.

We constructed an approximated convex hull that encloses a majority of the better performing points in the three-dimensional environmental metric space as shown in Figure 5. From the diagram we can also see that there is an opportunity to trade off different types of environmental impact. In Table 4, we show one specific generated concrete formula that is nearest neighbor to one of the extremal points used to construct the convex hull, for strength of 30±1 MPa and 40±1 MPa respectively.

3.2 Extrapolative Generation

A key hallmark of creativity is novelty: the ability to generate ideas or artifacts that are beyond the training dataset. Given the concrete formulations in the training data do not meet our desired sustainability performance metrics, we need creativity. We aim for generative algorithms that are extrapolative rather than simply interpolative in the training data [7]. To see that our generated formulations are indeed novel, consider a visualization based on the isomap dimensionality reduction technique [38] in Figure 6 for Mix 4 from Table 5 (chosen arbitrarily) together with 38 mixes from the training data [41] that have similar 28-day compressive strength. The isomap spatial layout aims to preserve distances among points in the high-dimensional space as closely as possible in the low-dimensional space and is based on all ingredients {cement, slag, fly ash, water, superplasticizer, coarse aggregate, and fine aggregate}, whereas the embedded "pie charts as markers" only show the fractions of the first three ingredients {cement, slag, fly ash} for ease of understanding.

As can be observed, whether considering only the first three ingredients, or considering all of the ingredients, the new formulation (marked by an arrow) is distinct from the others and therefore novel. In this sense, this new formulation is pushing the boundary of human creativity. The design principle that the AI-generated formulations follow is to considerably decrease cement by replacing with other cementitious materials such as fly ash and slag. This is much beyond current worldwide practices that restrict cement replacement to 20–25% [39] and is therefore quite novel.

3.3 Visualization for Concrete Design

On top of the three-dimensional environmental impact design space that we mentioned earlier, we also color each data point based on the predicted strength of the corresponding formula. Figure 7 shows the strength spectrum of the newly generated concrete formulas plotted in the environmental impact space for each concrete curing time group. These plots could serve as a visualization tool for concrete designers to quickly select newly generated formulas that meet the design requirements.

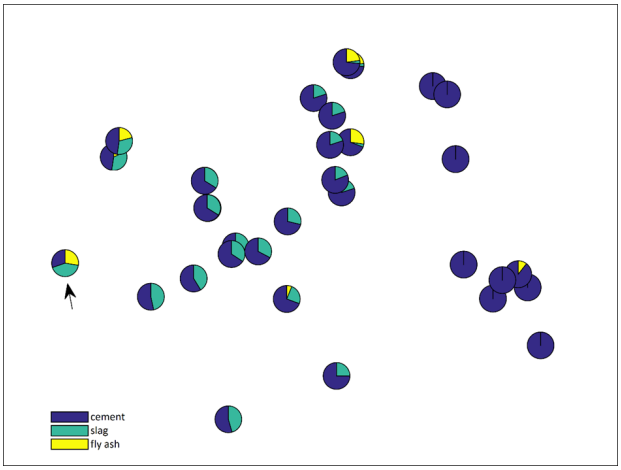

Figure 6: Isomap embedding of a novel AI-generated concrete formulation (marked by arrow), together with 38 formulations from the training data having similar 28-day strength. Placement is governed by all 7 ingredient dimensions but pie chart markers only show the cementitious ingredients. Note that isomap dimensions are unsupervised and must be interpreted in post hoc manner.

3.4 Strength-Conditioned Progression Generation

Attribute-conditioned generative progression has been considered by [34]. In the context of face image generation there, one of the attribute dimension values such as gender, facial expression, or hair color is modified by interpolating between the minimum and maximum attribute value, i.e. $x = [x_\alpha, x_{rest}]$, where $x_\alpha = (1-\alpha) \cdot x_{min} + \alpha \cdot x_{max}$. Indeed, one can visualize that the attribute of generated images change progressively with the change in conditioning attribute values.

To further demonstrate our concrete formulation generator can produce concrete designs with desired properties, we perform similar experiments. For the purpose of illustration, we limit our conditioning variables to strength and curing time of the concrete. We again generate 10,000 samples for each curing time group, by uniformly sampling from $[x_{min}, x_{max}]$. Figure 8 shows how well the predicted strength of generated formulations match with the desired strength given as conditioning variable during generation. The performance varies across different curing time groups. The RMSE is computed to evaluate the performance quantitatively. The better performing model should have the contour of the scattered dots to cover the diagonal line. The result shows that the generator seems to work the best for concrete curing time of 7 days.

4 LABORATORY EXPERIMENTS

Five concrete formulations were generated using the CVAE approach, aiming to minimize environmental impacts under given compressive strength targets. Since the nature of superplasticizer has changed since the training dataset was developed, human adjustment of superplasticizer proportion was made to improve rheology

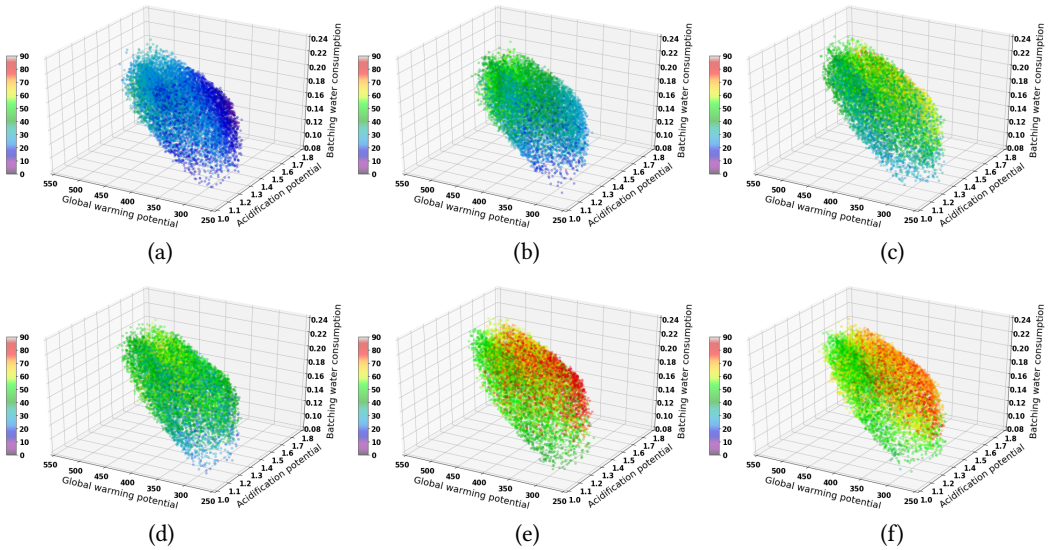

Figure 7: Strength spectrum of generated concrete formulas for different concrete curing time plotted in 3D environmental impacts space, where color indicates strength. (a) ≤ 3 days. (b) 7 days. (c) 14 days. (d) 28 days. (e) 56 days. (f) ≥90 days.

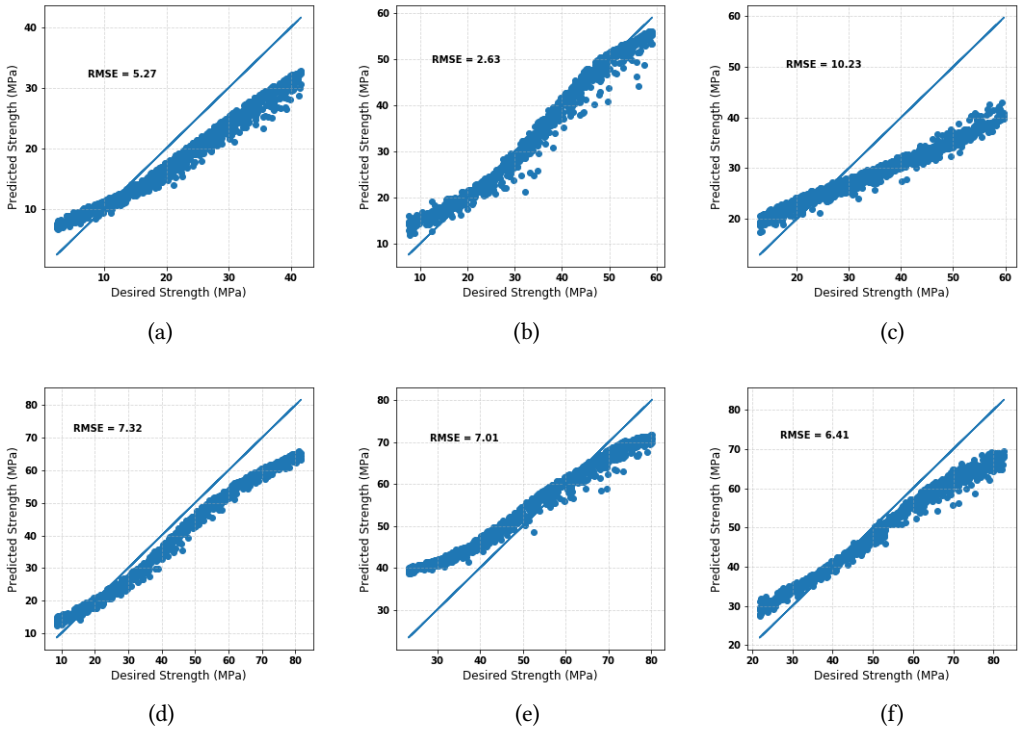

Figure 8: Strength-conditioned progression for different concrete curing times. (a) ≤ 3 days. (b) 7 days. (c) 14 days. (d) 28 days. (e) 56 days. (f) ≥90 days.

(and address this drift from historical data); in particular, all superplasticizer quantities were set to 1/4 of their AI-generated values. Proportions for the concrete mixes that were made for laboratory experiments are listed in Table 5.

To manufacture the concrete mixes, ordinary Portland cement (OPC), blast furnace slag, fly ash, and potable water were used for making the cement paste. Limestone and natural sand were used as coarse and fine aggregate, respectively. Both these aggregates meet the particle size distribution requirements from the ASTM standards. An industrial pan mixer was used for the mixing process.

4.1 Mixing Procedure

The properties of fresh concrete were investigated with concrete that was mixed in accordance with the ASTM C192 standard. First, the pan was wet with water, and the excess water was wiped out with a towel. Next, coarse aggregate, fine aggregate, and half of the water were mixed for 30 seconds. After 30 seconds, the mixer was stopped, and the cementitious materials were added to the pan. The mixer was restarted, and the remaining water was added carefully to the pan. The concrete was mixed for 3 minutes, rested for 1 minute in which the superplasticizer was added, and then mixed for an additional 2 minutes, as specified by the ASTM standards.

4.2 Slump Test

A slump test for fresh concrete was then conducted. The test was conducted in accordance with the ASTM C143 standard. First, the slump cone was placed on the mat and secured to the floor by stepping on the foothold of the cone. The cone was filled with fresh concrete up to 1/3 of its total volume and rodded 25 times with the 5/8-inch (15.8 mm) tamping rod. These steps were repeated two more times until the slump cone was entirely filled with concrete. The top of the slump cone was leveled with the trowel. Finally, the slump cone was carefully pulled directly upward in a period of 3 to 7 seconds, and the slump value was measured using a scale.

4.3 Preparation of Concrete Cylinders

Cylindrical specimens were prepared with a diameter of 4 inches (101.6 mm) and a length of 8 inches (203.2 mm) to be used for the compressive strength tests, according to ASTM C31. A total of 6 cylindrical specimens were made for each mix, three to be tested for compressive strength after 7 days, and the other three to be tested after 28 days. To prepare the cylinders, the cylindrical molds were first oiled to assist in demolding. Any excess oil was then wiped out with a towel to ensure there was only a thin layer of form oil. The first half of the cylinder was filled with concrete, rodded 25 times with a small tamping rod, then the outside of the mold was tapped 10 to 15 times to eliminate the air voids. The second half of the mold was then filled, tamped, and tapped in the same exact manner as the first half. The top of the mold was then finished with the trowel, and the cylinders were covered with a plastic sheet and left for 24 hours to harden. After hardening, the specimens were demolded and left at a temperature of 68°F (20°C) in the moist cure room.

Figure 9: Photographs of some concrete cylinders made according to novel formulations, undergoing laboratory testing for compressive strength.

4.4 Compressive Strength Tests

A Forney machine was used to conduct compressive strength tests on the concrete cylinders after 7 and 28 days of curing and in accordance with the ASTM C39 standard. To perform the test, the cylindrical specimen was first capped with sulfur-based capping compound on the top and bottom and placed on the center of the loading platen. The load was applied at a constant rate of 27,000 lb/min (12,240 kg/min) until the specimen failed, and the ultimate load was recorded. Also, the fracture pattern was noted.

Figure 9 shows some of the concrete cylinders, whereas Figure 10 shows compressive testing results. These results are further quantified in Table 6. Note that all five tested mixes exceed the target 28-day strength already at 7 days and that the 28-day measured strength was double or more than the target. Table 6 shows results from the slump test.

4.5 Global Warming Potential

Table 6 also shows the lifecycle analysis-based computations for the GWP, AP, and CBW environmental metrics. Table 7 explicitly demonstrates that these formulations roughly halve the global warming potential as compared to the average of similar 28-day compressive strength formulations. To make this comparison, we used the achieved strength numbers from Table 6 against similar achieved numbers from the UCI Machine Learning Repository [41].

Further comparison could be made against the industry standard given in [1], which gives regional performance benchmarks. Importantly, this provides an external comparison beyond the training dataset that was used in our work. For the Great Lakes region that the field test associated with this paper belongs to, this report gives the following benchmarks, against which our numbers compare favorably:

- 3000 psi: 281.33 kg CO_2/m^3
- 4000 psi: 334.87 kg CO_2/m^3

5 INDUSTRIAL TESTING AND FIELD TESTS

5.1 Industrial Testing and Adjustments for Logistical and Environmental Factors

The lab-tested formulations—in particular Mix 1 with 28-day compressive strength target of 4000 psi—were provided to concrete supplier Ozinga Ready Mix of Chicago, IL. Further industrial lab testing was done, and life cycle analysis was re-performed using

mix	cement [kg/m^3]	blast furnace slag [kg/m^3]	fly ash [kg/m^3]	water (used) [kg/m^3]	superplasticizer (used) [kg/m^3]	coarse aggregate [kg/m^3]	fine aggregate [kg/m^3]
1	131.46	201.21	119.67	180.70	1.1	950.72	780.48
2	128.59	197.46	124.24	184.31	1.0	954.48	787.47
3	134.89	182.74	113.78	179.43	0.9	953.22	785.28
4	132.25	184.37	119.74	181.03	1.8	954.10	786.55
5	129.02	210.60	122.80	184.63	1.0	953.50	780.11

Table 5: Concrete mixes that underwent laboratory testing.

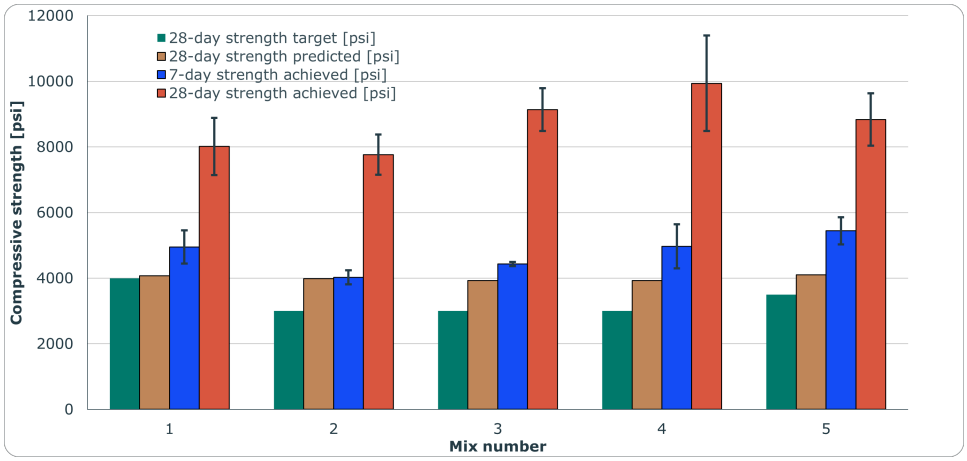

Figure 10: Experimental compressive strength results for five AI-generated concrete mixes. Error bars correspond to three cylinders each.

mix	28 day strength target [psi]	7 day strength [psi]	28 day strength [psi]	slump [in]	global warming potential [kg CO_2 eq./m^3]	acidification potential [kg SO_2 eq./m^3]	batch water [m^3]
1	4000	4949 ± 507	8013 ± 871	5 1/2	154.111	0.534	0.181
2	3000	4025 ± 217	7764 ± 610	3 1/4	152.284	0.530	0.183
3	3000	4431 ± 66	9136 ± 656	6 3/4	157.294	0.547	0.180
4	3000	4967 ± 673	9938 ± 1453	5 3/4	155.157	0.549	0.167
5	3500	5443 ± 417	8836 ± 801	3 3/4	152.149	0.524	0.169

Table 6: Laboratory testing and lifecycle analysis results.

mix	1	2	3	4	5
Estimated GWP (kg CO_2 eq./m^3)	154.11	152.28	157.29	155.16	152.15
Average of industry standard of similar 28-day compressive strength (kg CO_2 eq./m^3)	282.36	280.31	318.75	302.45	279.78

Table 7: Comparing global warming potential of AI-generated mixes to average of industry standard mixes of similar 28-day compressive strength.

the Athena/NRMCA EPD tool,[2] which is an alternative to the CSI EPD tool used by them. Under this life cycle analysis, Mix 1 has GWP of 171 kg CO_2/m^3.

[2]https://www.nrmca.org/wp-content/uploads/2020/02/AMSI_EPDtool.pdf

Changes were made to the original formula to take two factors into account:

(1) Logistical considerations, for example the availability of fly ash and slag in large quantities.

(2) Potentially cold weather conditions at the Meta Data Center construction site, which will reduce curing speed, especially with fly ash and slag as cementitious material [30, 36].

Since Mix 1 contained significant amounts of fly ash and slag as cementitious material, the formula was adjusted by Ozinga to address these two factors (with the expectation that field tests would be under cold conditions), as given in Table 8, for two separate field tests. Although the adjusted formulation had less slag and fly ash than the AI-derived formula, the GWP remained very close. The GWP of the revised Ozinga formula under the Athena life cycle analysis was 174 kg CO_2/m^3.

5.2 Deployment in Field Tests

After internal testing and adjustments by Ozinga, the formulations were deployed in two separate pours of slab-on-grade applications. A slab-on-grade is a type of shallow foundation in which a concrete slab rests directly on the ground below it. A slab-on-grade foundation usually consists of a thin layer of concrete across the entire area of the foundation with thickened footings at the edges.

In Pour 1, the concrete was used for an office building used for the construction crew. This slab-on-grade foundation of size 40,000 square feet, see Figure 11, was done in March 2021 during cold days, with temperatures in (42°F–52°F). In Pour 2, the concrete was used for a guard house. This slab-on-grade foundation of size 450 square feet, see Figure 12, was done in June 2021 during hot days, with temperatures in (80°F–98°F).

The quantitative results from these field deployments are given in Tables 9–11. As can be observed, all specifications were met, despite the variability in weather and other environmental conditions. The only concern raised by the contractor was the slow curing time of Pour 1 in cold weather, in missing the unspecified, but still desirable, 5-day strength at 75% of 28-day target. Early strength performance at 1, 3, or 5-day may be needed specifically for slab-on-grade applications to allow further activities such as saw cutting and other finishing. The qualitative results, such as visual aesthetics, were also acceptable.

(a)

(b)

Figure 11: Pictures of Pour 1 on day of deployment. (a) first half. (b) second half.

6 CONCLUSION

We have demonstrated end-to-end accelerated design and deployment of low-carbon concrete, based on a CVAE model that discovers new concrete formulations with reduced environmental impact, including embodied carbon, without sacrificing strength. The resulting formulations were then further optimized based on environmental and logistical constraints.

However, there is still room for improving our model and broader approach. Our data contains both continuous and categorical values, but CVAE may not be the best for capturing such mixed categorical and continuous data. The VAE-ROC model proposed by [37] is said to be better at handling mixed data. We hope by modifying the CVAE model in line with specifics of the VAE-ROC, the generator would synthesize more realistic concrete designs and achieve better performance in attribute-conditioned generation. Faster curing of concrete formulations in cold weather is also desired as an additional property. This includes 1, 3, and 5-day curing performance which are important for mission-critical applications such as slab-on-grade.

Finally, as we have noted, the availability of cement replacement materials such as fly ash and slag vary with location and time. Therefore there is an opportunity to directly optimize for such logistical considerations. Moreover, predicted weather conditions can be taken into account. Further, there is need to identify and discover novel materials that could be used in addition to, or in replacement of, such cementitious materials.

ACKNOWLEDGMENTS

Discussions with Jeremy Gregory, Randolph Kirchain, and Joana Maria are very much appreciated.

This work was supported in part by an unrestricted gift from Meta. XG, HY, and LRV were also supported in part by the IBM-Illinois Center for Cognitive Computing Systems Research (C3SR), a research collaboration as part of the IBM AI Horizons Network.

REFERENCES

[1] 2019. *A Cradle-to-Gate Life Cycle Assessment of Ready-Mixed Concrete Manufactured by NRMCA Members.* Technical Report. Athena Sustainable Materials Institute. Version 3.

	Pour 1	Pour 2	Mix 1
Ingredient	Amount per cubic yard	Amount per cubic yard	Amount per cubic yard
Cement (ASTM C 150 Type I/II)	215 lb	215 lb	211.6 lb
Fly Ash (ASTM C 618 Class C)	100 lb	100 lb	339.2 lb
Slag (ASTM C 989)	180 lb	180 lb	201.7 lb
Water – Potable	25.3 gal	24.0 gal	36.5 gal
Coarse Aggregate (ASTM C 33)	1930 lb	1920 lb	1602.5 lb
Fine Aggregate (ASTM C 33)	1500 lb	1485 lb	1315.5 lb
Water-Cementitious Material Ratio	0.43	0.40	0.40

Table 8: Compositions of two field-deployed concrete mixes. Note that the mixes also included CarbonCure, Accelerator (ASTM C 494 Type C), and Water Reducer (ASTM C 494 Type F) whose admixture dosage rates varied based on concrete temperature, ambient temperature, haul time, etc.

Performance Metric	Specification	Result	Note
3-day strength	—	1936 psi	no spec, but better than expected
5-day strength	—	2650 psi	no spec, but less than expected
7-day strength	3000 psi	3510 psi	passed spec
28-day strength	4000 psi	6970 psi	passed spec
Slump	5–7 in.	6.42 in.	passed spec
Air content	0–3%	1.3%	passed spec
Shrinkage	< 0.032% at 21 days (ASTM C157)	0.022%	passed spec

Table 9: Quantitative results from first half of Pour 1: compressive strengths are average of seven cylinders; slump and air content are averages of three samples.

Performance Metric	Specification	Result	Note
3-day strength	—	1620 psi	no spec, but better than expected
5-day strength	—	2667 psi	no spec, but less than expected
7-day strength	3000 psi	3540 psi	passed spec
28-day strength	4000 psi	7592 psi	passed spec
Slump	5–7 in.	6.25 in.	passed spec
Air content	0–3%	1.6%	passed spec
Shrinkage	< 0.032% at 21 days (ASTM C157)	0.022%	passed spec

Table 10: Quantitative results from second half of Pour 1: compressive strengths are average of seven cylinders; slump and air content are averages of three samples.

Performance Metric	Specification	Result	Note
3-day strength	—	1960 psi	no spec, but better than expected
7-day strength	—	2940 psi	no spec, but as expected
28-day strength	4000 psi	6260 psi	passed spec
Slump	5–7 in.	5.50 in.	passed spec
Air content	0–3%	2.2%	passed spec
Shrinkage	< 0.032% at 21 days (ASTM C157)	0.033%	passed spec (okay within 15%)

Table 11: Quantitative results from Pour 2: compressive strengths are average of seven cylinders; slump and air content are averages of three samples.

[2] 2021. Concrete needs to lose its colossal carbon footprint. *Nature* 597, 7878 (Sept. 2021), 593–594. https://doi.org/10.1038/d41586-021-02612-5

[3] Massimo Aria, Corrado Cuccurullo, and Agostino Gnasso. 2021. A comparison among interpretative proposals for Random Forests. *Machine Learning with Applications* 6 (Dec. 2021), 100094. https://doi.org/10.1016/j.mlwa.2021.100094

[4] Wassim Ben Chaabene, Majdi Flah, and Moncef L. Nehdi. 2020. Machine learning prediction of mechanical properties of concrete: Critical review. *Construction and Building Materials* 260 (Nov. 2020), 119889. https://doi.org/10.1016/j.conbuildmat.2020.119889

[5] Robert Courland. 2011. *Concrete Planet: The Strange and Fascinating Story of the World's Most Common Man-Made Material.* Prometheus.

(a)

(b)

Figure 12: Pictures of Pour 2. (a) the foundation. (b) the saw cut made into the foundation.

[6] F. Dabbaghi, A. Tanhadoust, M. L. Nehdi, S. Nasrollahpour, M. Dehestani, and H. Yousefpour. 2021. Life cycle assessment multi-objective optimization and deep belief network model for sustainable lightweight aggregate concrete. *Journal of Cleaner Production* 318 (Oct. 2021), 128554. https://doi.org/10.1016/j.jclepro.2021.128554

[7] Payel Das and Lav R. Varshney. 2022. Explaining AI Generation and Creativity. *IEEE Signal Processing Magazine* (2022). to appear.

[8] M. A. DeRousseau, J. R. Kasprzyk, and W. V. Srubar, III. 2018. Computational design optimization of concrete mixtures: A review. *Cement and Concrete Research* 109 (July 2018), 42–53. https://doi.org/10.1016/j.cemconres.2018.04.007

[9] Carl Doersch. 2016. Tutorial on Variational Autoencoders. arXiv:1606.05908 [stat.ML].

[10] Leah D. Ellis, Andres F. Badel, Miki L. Chiang, Richard J.-Y. Park, and Yet-Ming Chiang. 2020. Toward electrochemical synthesis of cement—An electrolyzer-based process for decarbonating $CaCO_3$ while producing useful gas streams. *Proceedings of the National Academy of Sciences of the United States of America* 117, 23 (June 2020), 12584–12591. https://doi.org/10.1073/pnas.1821673116

[11] Kathryn Furlong. 2021. Geographies of infrastructure II: Concrete, cloud and layered (in)visibilities. *Progress in Human Geography* 45, 1 (Feb. 2021), 190–198. https://doi.org/10.1177/0309132520923098

[12] Jim Gao. 2014. *Machine Learning Applications for Data Center Optimization.* Technical Report. Google.

[13] Xiou Ge, Richard T. Goodwin, Jeremy R. Gregory, Randolph E. Kirchain, Joana Maria, and Lav R. Varshney. 2019. Accelerated Discovery of Sustainable Building Materials. In *AAAI 2019 Spring Symposium, Towards AI for Collaborative Open Science (TACOS).*

[14] Abhay Goyal, Ivan Palaia, Katerina Ioannidou, Franz-Josef Ulm, Henri van Damme, Roland J.-M. Pellenq, Emmanuel Trizac, and Emanuela Del Gado. 2021. The physics of cement cohesion. *Science Advances* 7, 32 (Aug. 2021), eabg5882. https://doi.org/10.1126/sciadv.abg5882

[15] Rafael Gómez-Bombarelli, Jennifer N. Wei, David Duvenaud, José Miguel Hernández-Lobato, Benjamín Sánchez-Lengeling, Dennis Sheberla, Jorge Aguilera-Iparraguirre, Timothy D. Hirzel, Ryan P. Adams, and Alán Aspuru-Guzik. 2018. Automatic Chemical Design Using a Data-Driven Continuous Representation of Molecules. *ACS Central Science* 4, 2 (Feb. 2018), 268–276. https://doi.org/10.1021/acscentsci.7b00572

[16] Taihao Han, Ashfia Siddique, Kamal Khayat, Jie Huang, and Aditya Kumar. 2020. An ensemble machine learning approach for prediction and optimization of modulus of elasticity of recycled aggregate concrete. *Construction and Building Materials* 244 (May 2020), 118271. https://doi.org/10.1016/j.conbuildmat.2020.118271

[17] Per Jahren and Tongbo Sui. 2018. *History of Concrete: A Very Old and Modern Material.* World Scientific / Chemical Industry Press.

[18] Anubhav Jain, Shyue Ping Ong, Geoffroy Hautier, Wei Chen, William Davidson Richards, Stephen Dacek, Shreyas Cholia, Dan Gunter, David Skinner, Gerbrand Ceder, and Kristin A. Persson. 2013. The Materials Project: A materials genome approach to accelerating materials innovation. *APL Materials* 1 (2013), 011002. https://doi.org/10.1063/1.4812323

[19] Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL: A Conditional Transformer Language Model for Controllable Generation. arXiv:1909.05858 [cs.CL].

[20] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning Representations (ICLR).*

[21] Diederik P. Kingma, Danilo J. Rezende, Shakir Mohamed, and Max Welling. 2014. Semi-supervised Learning with Deep Generative Models. In *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Eds.). Curran Associates, Inc., 3581–3589.

[22] Diederik P. Kingma and Max Welling. 2019. An Introduction to Variational Autoencoders. *Foundations and Trends in Machine Learning* 12, 4 (Nov. 2019), 307–392. https://doi.org/10.1561/2200000056

[23] Julie Kriegh, Chris Magwood, and Wil Srubar, III. 2021. *Carbon-Storing Materials.* Summary Report. Carbon Leadership Forum.

[24] Nina Kshetry and Lav R. Varshney. 2019. Safety in the Face of Unknown Unknowns: Algorithm Fusion in Data-Driven Engineering Systems. In *Proceedings of the 2019 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP).* 8162–8166. https://doi.org/10.1109/ICASSP.2019.8683392

[25] Turab Lookman, Stephan Eidenbenz, Frank Alexander, and Cris Barnes. 2018. *Materials Discovery and Design: By Means of Data Science and Optimal Learning.* Springer Nature. https://doi.org/10.1007/978-3-319-99465-9

[26] Soroush Mahjoubi, Rojyar Barhemat, Pengwei Guo, Weina Meng, and Yi Bao. 2021. Prediction and multi-objective optimization of mechanical, economical, and environmental properties for strain-hardening cementitious composites (SHCC) based on automated machine learning and metaheuristic algorithms. *Journal of Cleaner Production* 329 (Dec. 2021), 129665. https://doi.org/10.1016/j.jclepro.2021.129665

[27] Eric Masanet, Arman Shehabi, Nuoa Lei, Sarah Smith, and Jonathan Koomey. 2020. Recalibrating global data center energy-use estimates. *Science* 367, 6481 (Feb. 2020), 984–986. https://doi.org/10.1126/science.aba3758

[28] Alexander Meinke and Matthias Hein. 2020. Towards neural networks that provably know when they don't know. In *Proceedings of the International Conference on Learning Representations (ICLR).*

[29] Reese Palley. 2010. *Concrete: A Seven-Thousand-Year History.* Quantuck Lane Press.

[30] Walter H. Price. 1951. *Factors Influencing Concrete Strength.* Technical Report 2. American Concrete Institute. 417–432 pages.

[31] Ladislav Rampášek, Daniel Hidru, Petr Smirnov, Benjamin Haibe-Kains, and Anna Goldenberg. 2019. Dr.VAE: improving drug response prediction via modeling of drug perturbation effects. *Bioinformatics* 35, 19 (Oct. 2019), 3743–3751. https://doi.org/10.1093/bioinformatics/btz158

[32] Jiaolong Ren, Hongbo Zhao, Lin Zhang, Zedong Zhao, Yinshan Xu, Yuanyuan Cheng, Meng Wang, Jingchun Chen, and Jian Wang. 2022. Design optimization of cement grouting material based on adaptive boosting algorithm and simplicial homology global optimization. *Journal of Building Engineering* 49 (May 2022), 104049. https://doi.org/10.1016/j.jobe.2022.104049

[33] David Rolnick, Priya L. Donti, Lynn H. Kaack, Kelly Kochanski, Alexandre Lacoste, Kris Sankaran, Andrew Slavin Ross, Nikola Milojevic-Dupont, Natasha Jaques, Anna Waldman-Brown, Alexandra Luccioni, Tegan Maharaj, Evan D. Sherwin, S. Karthik Mukkavilli, Konrad P. Kording, Carla Gomes, Andrew Y. Ng, Demis Hassabis, John C. Platt, Felix Creutzig, Jennifer Chayes, and Yoshua Bengio. 2019. Tackling Climate Change with Machine Learning. arXiv:1906.05433 [cs.CY].

[34] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. 2015. Learning Structured Output Representation using Deep Conditional Generative Models. In *Advances in Neural Information Processing Systems 28*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett (Eds.). Curran Associates, Inc., 3483–3491.

[35] Y. Song, B. Ouyang, J. Chen, X. Wang, K. Wang, S. Zhang, Y. Chen, G. Sant, and M. Bauchy. 2022. Decarbonizing concrete with artificial intelligence. In *Computational Modelling of Concrete and Concrete Structures.* CRC Press, 168–176. https://doi.org/10.1201/9781003316404-21

[36] Marios Soutsosa and Fragkoulis Kanavarisb. 2018. The modified nurse-saul (MNS) maturity function for improved strength estimates at elevated curing temperatures. *Case Studies in Construction Materials* 9 (Dec. 2018), e00206. https:

//doi.org/10.1016/j.cscm.2018.e00206

[37] Suwon Suh and Seungjin Choi. 2016. Gaussian Copula Variational Autoencoders for Mixed Data. arXiv:1604.04960 [stat.ML].

[38] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. 2000. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science* 290, 5500 (Dec. 2000), 2319–2323. https://doi.org/10.1126/science.290.5500.2319

[39] UN Environment, Karen L. Scrivener, Vanderley M. John, and Ellis M. Gartner. 2018. Eco-efficient cements: Potential economically viable solutions for a low-CO_2 cement-based materials industry. *Cement and Concrete Research* 114 (Dec. 2018), 2–26. https://doi.org/10.1016/j.cemconres.2018.03.015

[40] Xinchen Yan, Jimei Yang, Kihyuk Sohn, and Honglak Lee. 2016. Attribute2Image: Conditional Image Generation from Visual Attributes. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 776–791. https://doi.org/10.1007/978-3-319-46493-0_47

[41] I.-C. Yeh. 1998. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research* 28, 12 (Dec. 1998), 1797–1808. https://doi.org/10.1016/S0008-8846(98)00165-3

[42] Benjamin A. Young, Alex Hall, Laurent Pilon, Puneet Gupta, and Gaurav Sant. 2019. Can the compressive strength of concrete be estimated from knowledge of the mixture proportions?: New insights from statistical analysis and machine learning methods. *Cement and Concrete Research* 115 (Jan. 2019), 379–388. https://doi.org/10.1016/j.cemconres.2018.09.006

[43] Junfei Zhang, Yimiao Huang, Guowei Ma, and Brett Nener. 2021. Mixture optimization for environmental, economical and mechanical objectives in silica fume concrete: A novel frame-work based on machine learning and a new metaheuristic algorithm. *Resources, Conservation and Recycling* 167 (April 2021), 105395. https://doi.org/10.1016/j.resconrec.2021.105395

[44] Junfei Zhang, Yimiao Huang, Yuhang Wang, and Guowei Ma. 2020. Multi-objective optimization of concrete mixture proportions using machine learning and metaheuristic algorithms. *Construction and Building Materials* 253 (Aug. 2020), 119208. https://doi.org/10.1016/j.conbuildmat.2020.119208